arXiv:1611.07214v1 [cs.IT] 22 Nov 2016

COMPARING ENTROPY RATES ON FINITE AND INFINITE ROOTED TREES WITH LENGTH FUNCTIONS

THOMAS HIRSCHLER, WOLFGANG WOESS

Abstract. We consider stochastic processes with (or without) memory whose evolution is encoded by a finite or infinite rooted tree. The main goal is to compare the entropy rates of a given base process and a second one, to be considered as a perturbation of the former. The processes are described by probability measures on the boundary of the given tree, and by corresponding forward transition probabilities at the inner nodes. The comparison is in terms of Kullback-Leibler divergence. We elaborate and extend ideas and results of Böcherer and Amjad. Our extensions involve length functions on the edges of the tree as well as nodes with countably many successors. In the last part, we consider trees with infinite geodesic rays and random perturbations of a given process.

1. Introduction

The relation between probability distributions on the set of leaves of a finite, rooted tree and forward transition probabilities at the internal nodes in the context of the analysis of variable length codes goes back to Massey [8], [9]. In [1], Böcherer and Amjad take up this setting. They derive a bound on the difference of the entropies between two such distributions, where the second one is "memoryless" in terms of an encoding of the tree in terms of strings over a finite alphapet. The bound is in terms of the Kullback-Leibler divergence of the first with respect to the second distribution, and of the lenghts (heights) of the leaves.

In this note, we generalise these inequalities, thereby also streamlining the mathematical substance. Our generalisation consists in (1) comparing the entropies of two arbitrary distributions on the leaves in terms of the Kullback-Leibler divergence, (2) involving more general length functions than those where each edge of the tree has length 1, and (3) allowing internal nodes with a countable infinity of forward neighbours. Finally, (4), we also consider trees which have infinite geodesic rays, corresponding to boundary points at infinity.

In the first main parts of this paper, the boundary ∂T of our tree T with root o consists of leaves (vertices with no forward neighbour) and is at most countably infinite. It is equipped with two probability measures $\mathbb{P}$ and $\mathbb{Q}$, where $\mathbb{Q}$ plays the role of a fixed reference measure and $\mathbb{P}$ should be thought of as a perturbation of $\mathbb{Q}$ which corresponds to some experimental situation.

The internal vertices (nodes) x of T are equipped with a function $\ell(x) > 0$ with the interpretation that each outgoing edge at x has length $\ell(x)$. The absolute value $|x|_\ell$ of any node is then the sum of the lengths of the edges on the unique geodesic path from o to x. The habitual choice is $\ell = \ell_\sharp \equiv 1$. Thus we have the expected length $\ell(\mathbb{P}) = \sum_{v \in \partial T} |v|_\ell \, \mathbb{P}(v)$, and analogously $\ell(\mathbb{Q})$. Let $\mathsf{H}(\mathbb{P})$ and $\mathsf{H}(\mathbb{Q})$ be the entropies of the two measures, and $\mathsf{D}(\mathbb{P}\|\mathbb{Q})$ the Kullback-Leibler divergence of $\mathbb{P}$ with respect to $\mathbb{Q}$.

Key words and phrases. entropy rate, Kullback-Leibler divergence, tree, boundary, leaf-node interchange, stochastic process.

The second author was partially supported by Austrian Science Fund project FWF P-24028.

The aim is to estimate the difference between $\mathsf{H}(\mathbb{P})/\ell(\mathbb{P})$ and $\mathsf{H}(\mathbb{Q})/\ell(\mathbb{Q})$ in terms of $\mathsf{D}(\mathbb{P}\|\mathbb{Q})/\ell(\mathbb{P})$.

The probability measure $\mathbb{P}$ gives rise to a forward Markov chain on the internal nodes of T, whose transition probabilities are $p(y|x) = \mathbb{P}(\partial T_y)/\mathbb{P}(\partial T_x)$, when y is a forward neighbour of x, and where ∂T_x denotes the set of leaves v for which x lies on the geodesic path from o to v. The analogous transition probabilites associated with $\mathbb{Q}$ are denoted $q(y|x)$. A basic tool for achieving the desired estmate is the "Leaf-Average Node-Sum Interchange Theorem" (LANSIT) of Rueppel and Massey [9]. It relates $\mathsf{H}(\mathbb{P})/\ell(\mathbb{P})$ with the entropies $\mathsf{H}\big(p(\cdot|x)\big)$ and $\mathsf{D}(\mathbb{P}\|\mathbb{Q})/\ell(\mathbb{P})$ with the divergences $\mathsf{D}\big(p(\cdot|x)\|q(\cdot|x)\big)$.

The main concern of [1] regards the situation when the tree is finite, $\ell = \ell_\sharp$, and all probability measures $q(\cdot|x) = q(\cdot)$ coincide under a suitable labelling of the sets of forward neighbours of the internal nodes x. In this case, the LANSIT yields $\mathsf{H}(\mathbb{P})/\ell_\sharp(\mathbb{Q}) = \mathsf{H}(q)$, and the estimate of [1] can be read as

$$\left|\frac{\mathsf{H}(\mathbb{P})}{\ell_\sharp(\mathbb{P})} - \mathsf{H}(q)\right| < M\varphi(\delta) + \frac{C}{\delta}\sqrt{\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell_\sharp(\mathbb{P})}}\,, \tag{1.1}$$

where δ is an arbitrarily small positive constant, M is the (constant) number of forward neighbours of the internal nodes, and $C \approx 1.25$.

We start in §2 by explaining our notation, and by reviewing the proof of the LANSIT in presence of a general length function as above. As mentioned, our tree is allowed to have nodes with infinitely many forward neighbours. Thus, our entropies and divergences can be infinite series, and we require the latter to be absolutely convergent. Furthermore, we have local entropies $\mathsf{H}\big(p(\cdot|x)\big)$ and $\mathsf{H}\big(q(\cdot|x)\big)$ at the possibly infinitely many internal nodes, and we require the defining series to be uniformly summable. See §3, where this is discussed, and further basics are presented.

Theorem 4.1 in §4 contains our generalisation of (1.1). We consider arbitary probability measures $\mathbb{P}$ and $\mathbb{Q}$ that are supported by ∂T and satisfy the mentioned summability conditions. Consequently, our estimate of $|\mathsf{H}(\mathbb{P})/\ell(\mathbb{P}) - \mathsf{H}(\mathbb{Q})/\ell(\mathbb{Q})|$ is a sum of three terms, where the first can be made arbitrarily small, the central one comprises $\sqrt{\mathsf{D}(\mathbb{P}\|\mathbb{Q})/\ell(\mathbb{P})}$, and the last one is a multiple of the variational difference between the node-average probability measures induced by $\mathbb{P}$ and $\mathbb{Q}$ on the set of internal vertices. The pre-factor of this term vanishes precisely when $\mathsf{H}\big(q(\cdot|x)\big)/\ell(x)$ is constant (or, in the last part, asymptotically constant) on the internal nodes. Our assumptions are as general as possible within reasonable feasibility.

After a discussion of various aspects of our estimate (§5), in §6 we turn our attention to trees which have infinite geodesic paths. In this case, the boundary of the tree consists of the space of "ends" (limit points of infinite paths) and possibly also leaves. Applying Theorem 4.1 yields estimates for the entropy rate and variants thereof, see Theorem 6.4. In particular, coming back to the natural length function $\ell_\sharp$, if $\mathbb{Q}$ is the probability distribution on the trajectory space of a stochastic process with asymptotically constant, finite forward entropies $\mathsf{H}\big(q(\cdot|x)\big)$ on the associated infinite tree and $\mathbb{P}$ is the probability governing a perturbation of that process, then finiteness of $\mathsf{D}(\mathbb{P}|\mathbb{Q})$ implies that the perturbed process has the same entropy rate as the non-perturbed one. Note again that we allow the state space of these processes to be countably infinite, as long as $\mathsf{H}\big(q(\cdot|x)\big)$ is finite. Moreover, this also holds when $\mathsf{D}(\mathbb{P}_n|\mathbb{Q}_n)/n \to 0$, where $\mathbb{P}_n$ and $\mathbb{Q}_n$ are the distributions

of the respective processes up to time n. A final application is given in Teorem 6.10, where we consider *random* perturbations of a given stochastic process.

We remark that in a somewhat different vein, there has been work on the statistics of ergodic Markov chains. Ciuperca, Girardin and Lhote [2], [3], as well as Yari and Nikooravesh [13], consider the approximation of the entropy rate via the natural estimator (in terms of relative frequencies of transitions between two states, resp. visits in one state) of the chain at time n, as $n \to \infty$. Constantness of local entropies of the chain is not needed here. Note that this is not comparison with one perturbed chain, but with a sequence of random perturbations: the estimated transition probabilites and invariant distribution are updated at each time step. In [2], the state space is finite, and in [13], at each state there are only finitely many positive outgoing transition probabilities. Both correspond to situations where the corresponding infinite tree is locally finite. In [3], this finite range assumption is replaced by a quasi-power property.

2. Setting, notation, and preliminary facts

A. Trees, leaves, and lengths

A *rooted tree* is a connected graph T without circles, with one vertex o desginated as the root. We shall tacitly identify T with the set of vertices of the tree, and write $E(T)$ for the set of edges, if needed.

A *geodesic* or *geodesic path* in T is a sequence $\pi = [x_0\,, x_1\,, \dots, x_n]$ of successive neighbours in T such that all x_i are distinct. For any pair of vertices x, y there is precisey one geodesic $\pi(x, y)$ from x to y. Analogously, an infinite geodesic has the form $\pi = [x_0\,, x_1\,, x_2\,, \dots]$ where x_{k-1} and x_k are neighbours for each $k \in \mathbb{N}$.

While we do not necessarily assume T to be finite, *we shall mostly assume that it contains no infinite geodesic,* with the exception of the last §6.

Every vertex $x \in T\setminus\{o\}$ has a unique *predecessor* x^-, the neighbour of x on the geodesic $\pi(o, x)$. In this case, we say that x is a *successor* of x^-, and write $N(x)$ for the set of sucessors of $x \in T$. The *forward degree* $\deg^+(x)$ of $x \in T$ is the number of successors of x. In our setting, it will be natural to assume that $\deg^+(x) \neq 1$ for all $x \in T$. The set of *leaves* or *boundary* ∂T of T consists of all vertices which have no successor. By our assumption, it is non-empty, and every vertex lies on a geodesic from o to some leaf. The *interior* of T is $\operatorname{int} T = T \setminus \partial T$.

We can describe each edge of T as $e = [x^-, x]$, where $x \in T \setminus \{o\}$. We now assign a *length* $\ell(e) > 0$ to each edge e.

- *We restrict attention to length functions where $\ell(e)$ depends only on x^- for $e = [x^-, x]$.*

That is, at each vertex the outgoing edges have the same length, and for e as above we can write $\ell(e) = \ell(x^-)$. The length of a path is defined as the sum of the lengths of its edges. It induces a metric on T. For a vertex x, its length (distance to o) $|x|_\ell$ is the length of $\pi(o, x)$.

The standard length function is induced by $\ell_\sharp(e) = 1$. The associated path length is the number of edges. For $x \in T$, we then just write $|x|$ for the resulting length of $\pi(o, x)$. This is the ordinary *height* of x.

B. Forward Markov chains and leaf distributions

On a rooted tree, a *forward Markov chain* is given by transition probabilities

$$p(y|x)\,, \quad x \in \mathsf{int}\,T\,,\ y \in N(x)\,, \quad \text{with} \sum_{y \in N(x)} p(y|x) = 1\,.$$

We can interpret this via a random process: a particle starts at o. If at some time it is at a vertex $x \in \mathsf{int}\,T$ then the particle moves to a random successor of x chosen according to the probability distribution $p(\cdot|x)$. When the particle reaches a leaf, it stops there (so that we might add the trivial transition probabilites $p(v|v) = 1$ for $v \in \partial T$).

We now define a measure (function) on ∂T, as follows.

$$\mathbb{P}(v) = \prod_{x \in \pi(o,v)\setminus\{o\}} p(x|x^-)\,, \quad v \in \partial T\,. \tag{2.1}$$

This is the *hitting distribution* on that set: for a leaf v, we have that $\mathbb{P}(v)$ is the probability that our forward Markov chain terminates at v. Conversely, given any probability measure $\mathbb{P}$ on ∂T, it induces a forward Markov chain as follows.

(2.2) Definition. For $x \in T$, the *cone* or *branch* at x is the subtree with root x given by

$$T_x = \{y \in T : x \in \pi(o,y)\}$$

Its boundary ∂T_x consists of all $v \in \partial T$ with $x \in \pi(o,v)$. Then we define

$$p(x|x^-) = \mathbb{P}(\partial T_x)/\mathbb{P}(\partial T_{x^-})\,,$$

where of course $\mathbb{P}(\partial T_x) = \sum_{v \in \partial T_x} \mathbb{P}(v)$. Thus, we have a one-to-one correspondence between probability distributions on ∂T and forward Markov chains on the rooted tree T. See [9] for an introduction to these ideas. More generally, we can formulate the following.

(2.3) Definition. A *cross section* or *cut* is a subset $S \subset T$ such that for every leaf $v \in \partial T$, the geodesic $\pi(o,v)$ meets S in precisely one point.

We write T^S for the subtree with root o consisting of all vertices that lie on some geodesic $\pi(o,s)$ with $s \in S$.

We have of course $\partial T^S = S$ for any section. We define $\mathbb{P}_S$ on S in the same way as $\mathbb{P}$ was defined on ∂T in (2.1), so that $\mathbb{P}_S(x) = \mathbb{P}(\partial T_x)$. We get the hitting distribution of the forward Markov chain on S, a probability distribution.

Given our length function ℓ, the *expected length* of a section S with respect to $\mathbb{P}_S$ is

$$\ell(\mathbb{P}_S) = \sum_{x \in S} |x|_\ell\, \mathbb{P}_S(x)\,, \quad \text{in particular,} \quad \ell(\mathbb{P}) = \ell(\mathbb{P}_{\partial T})$$

We shall always assume that $\ell(\mathbb{P}) < \infty$. At last, given the probability distribution $\mathbb{P}$ on ∂T, we define new a new *(node average)* measure $\mu_{\mathbb{P}}$ on $\mathsf{int}\,T$ by

$$\mu_{\mathbb{P}}(x) = \frac{\ell(x)}{\ell(\mathbb{P})}\,\mathbb{P}(\partial T_x)\,. \tag{2.4}$$

It is a probability measure: we use the fact that $v \in \partial T_x \iff x \in \pi(o, v^-)$ for any leaf v and internal node x, whence

$$\sum_{x \in \mathsf{int}\, T} \ell(x)\, \mathbb{P}(\partial T_x) = \sum_{x \in \mathsf{int}\, T} \ell(x) \sum_{v \in \partial T_x} \mathbb{P}(v) = \sum_{v \in \partial T} \mathbb{P}(v) \sum_{x \in \pi(o, v^-)} \ell(x) = \ell(\mathbb{P})\,.$$

C. Gradient, Laplacian, and leaf-node interchange

(2.5) Definition. The *gradient* ∇ and *Laplacian* $\Delta = \Delta_{\mathbb{P}}$ on T are as follows. For a function $f : T \to \mathbb{R}$

$$\nabla f(y) = \frac{f(y) - f(y^-)}{\ell(y^-)}\,,\ y \in T \setminus \{o\}\,, \quad \text{and} \quad \Delta f(x) = \sum_{y \in N(x)} \nabla f(y)\, p(y|x)\,.$$

Some remarks are due: here, our notation differs from [1] as well as the old note [9], who write Δ for our ∇. In general, one should see ∇f as a function on $E(T)$, where f could be interpreted as a voltage function, and for $e = [y^-, y]$, the difference $f(y) - f(y^-)$ is the voltage drop, and thinking of $\ell(e)$ as the resistance of the edge, $\nabla f(e) = \big(f(y) - f(y^-)\big)/\ell(e)$ can be interpreted as Ohm's law. With our assumption, $\ell(e) = \ell(y^-)$, and we can adopt the above notation. Regarding the Laplacian

$$\Delta f(x) = \frac{1}{\ell(x)} \sum_{y \in N(x)} \big(f(y) - f(x)\big)\, p(y|x)\,,$$

as long as the foward degrees are finite, it is well defined for any function f. In presence of infinitely many successors, we require that the series converges absolutely:

$$|\Delta_{\mathbb{P}}| f(x) := \frac{1}{\ell(x)} \sum_{y \in N(x)} \big|f(y) - f(x)\big|\, p(y|x) < \infty \quad \text{for every} \ \ x \in \mathsf{int}\, T\,. \tag{2.6}$$

Our notation is compatible with the one of reversible Markov chain theory, see e.g. Woess [12, Chapter 4], although the present forward Markov chains are of course not reversible.

For a probability distribution ν on any finite or countable set $\mathcal{X}$, and any function $f : \mathcal{X} \to \mathbb{R}$, we write

$$\mathbb{E}_\nu(f) = \sum_{x \in \mathcal{X}} f(x)\, \nu(x)\,,$$

whenever f is ν-integrable, that is $\mathbb{E}_\nu(|f|) < \infty$. In this setting, the "Leaf-average node-sum interchange theorem" (LANSIT) of [9] is the following application of the discrete version of Fubini's theorem.

(2.7) Lemma. *Let $\mathbb{P}$ be a probability distribution on ∂T and $p(\cdot|\cdot)$ the associated forward Markov chain. Let $f : T \to \mathbb{R}$ be a function which satisfies $\mathbb{E}_{\mu_{\mathbb{P}}}\big(|\Delta_{\mathbb{P}}| f\big) < \infty$. Then f is $\mathbb{P}$-integrable on ∂T, and*

$$\frac{1}{\ell(\mathbb{P})}\, \mathbb{E}_{\mathbb{P}}\big(f - f(o)\big) = \mathbb{E}_{\mu_{\mathbb{P}}}\big(\Delta_{\mathbb{P}} f\big)$$

We remark here that in [9] – where $\ell = \ell_\sharp$ – it is suggested in the last lines of p. 347 that already the condition $\mathbb{E}_{\mu_{\mathbb{P}}}\big(|\Delta_{\mathbb{P}} f|\big) < \infty$ is sufficient for the conclusion of their Theorem 1. But this appears to be problematic, since it does not allow to separate in two parts the

sum as displayed in the middle of the proof on that page. For safety's sake, we present our proof.

Proof of Lemma 2.7. We start with the right hand side:

$$\begin{aligned}
\mathbb{E}_{\mu_{\mathbb{P}}}\big(\Delta_{\mathbb{P}} f\big) &\overset{(1)}{=} \frac{1}{\ell(\mathbb{P})} \sum_{x \in \mathsf{int}\, T} \sum_{y \in N(x)} \big(f(y) - f(x)\big) p(y|x) \mathbb{P}(\partial T_x) \\
&\overset{(2)}{=} \frac{1}{\ell(\mathbb{P})} \sum_{y \in T \setminus \{o\}} \big(f(y) - f(y^-)\big) \mathbb{P}(\partial T_y) \\
&\overset{(3)}{=} \frac{1}{\ell(\mathbb{P})} \sum_{y \in T \setminus \{o\}} \sum_{v \in \partial T_y} \big(f(y) - f(y^-)\big) \mathbb{P}(v) \\
&\overset{(4)}{=} \frac{1}{\ell(\mathbb{P})} \sum_{v \in \partial T} \sum_{y \in \pi(o,v) \setminus \{o\}} \big(f(y) - f(y^-)\big) \mathbb{P}(v) = \frac{1}{\ell(\mathbb{P})} \sum_{v \in \partial T} \big(f(y) - f(o)\big) \mathbb{P}(v)\,.
\end{aligned}$$

By assumption, the right hand side of (1) is absolutely convergent, so that the sum can be rearranged to get (2) and thus (3), which are also absolutely convergent. Thus, the discrete version of Fubini's theorem allows to exchange the order of summation, leading to (4). □

3. Entropies and tightness

A. Entropies and relative entropies

The *entropy* of our probability measure $\mathbb{P}$ on ∂T is

$$\mathsf{H}(\mathbb{P}) = \sum_{v \in \partial T} \mathbb{P}(v) \log_2 \frac{1}{\mathbb{P}(v)}\,.$$

When we have infinite forward degrees in T then ∂T is infinite. In this case, *we assume that* $\mathsf{H}(\mathbb{P}) < \infty$. At each inner node x of T, we have the outgoing probability distribution $p(\cdot|x)$ on the set of successors of x. Its entropy is defined analogously,

$$\mathsf{H}_p(x) = \mathsf{H}\big(p(\cdot|x)\big) = \sum_{y \in N(x)} p(y|x) \log_2 \frac{1}{p(y|x)}.$$

(3.1) Lemma. *Our assumption that* $\mathsf{H}(\mathbb{P}) < \infty$ *implies finiteness of* $\mathsf{H}_p(x)$ *for each* x*, and*

$$\frac{\mathsf{H}(\mathbb{P})}{\ell(\mathbb{P})} = \mathbb{E}_{\mu_{\mathbb{P}}}\left(\frac{\mathsf{H}_p(x)}{\ell(x)}\right).$$

On the right hand side, x is the variable with respect to which the expectation is taken.

Proof. As in [9], we set $f(x) = -\log_2 \mathbb{P}(\partial T_x)$. Then $\Delta_{\mathbb{P}} f \geq 0$, so that the proof of Lemma 2.7 goes through no matter whether the involved expectations are finite or not. Thus, the stated formula holds, and when $\mathsf{H}(\mathbb{P}) < \infty$, then also $\mathsf{H}_p(x)$ is finite for all x. □

Note in particular the following special choice of the length function.

$$\text{If} \quad \ell(x) = \mathsf{H}_p(x) \quad \text{then} \quad \ell(\mathbb{P}) = \mathsf{H}(\mathbb{P})\,. \tag{3.2}$$

Now let us take a second probability measure $\mathbb{Q}$ on ∂T, and denote the transition probabilities of the associated forward Markov chain by $q(y|x)$, when $y^- = x$. Then we

have the well-known *Kullback-Leibler divergence* or *relative entropy* of $\mathbb{P}$ with respect to $\mathbb{Q}$,

$$\mathsf{D}(\mathbb{P}\|\mathbb{Q}) = \sum_{v\in\partial T} \mathbb{P}(v) \log_2 \frac{\mathbb{P}(v)}{\mathbb{Q}(v)}\,.$$

For our purposes, it will always suffice to consider the non-degenerate situation when $\mathbb{P}(v)\,,\ \mathbb{Q}(v) > 0$ for each leaf v. In the case when ∂T is infinite, we assume again that the above series converges absolutely.

Similarly, we have the local version

$$\mathsf{D}_{p,q}(x) = \mathsf{D}\big(p(\cdot|x)\|q(\cdot|x)\big)\,,\quad x\in\mathsf{int}\,T\,.$$

We need again absolute convergence of its defining series. The following generalises the corresponding result of [1] to the case where T has vertices with infinite forward degree.

(3.3) Proposition. *If* $\mathsf{H}(\mathbb{P}) < \infty$ *and the defining series of* $\mathsf{D}(\mathbb{P}\|\mathbb{Q})$ *is absolutely convergent, then the same holds for each* $\mathsf{D}_{p,q}(x)$*, and*

$$\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell(\mathbb{P})} = \mathbb{E}_{\mu_{\mathbb{P}}}\!\left(\frac{\mathsf{D}_{p,q}(x)}{\ell(x)}\right).$$

Again, x is the variable with respect to which the expectation is taken on the right hand side.

Proof. We want to apply Lemma 2.7 to the function $f(x) = \log_2\big(\mathbb{P}(\partial T_x)/\mathbb{Q}(\partial T_x)\big)$. For this function, we compute

$$\begin{aligned}
\ell(\mathbb{P})\,\mathbb{E}_{\mu_{\mathbb{P}}}\big(|\Delta_{\mathbb{P}}|f\big) &= \sum_{x\in\mathsf{int}\,T} \mathbb{P}(\partial T_x) \sum_{y\in N(x)} p(y|x) \left|\log_2 \frac{p(y|x)}{q(y|x)}\right| \\
&= \sum_{y\in T\setminus\{o\}} \sum_{v\in\partial T_y} \mathbb{P}(v) \left|\log_2 \frac{p(y|x)}{q(y|x)}\right| \\
&\le \sum_{v\in\partial T} \mathbb{P}(v) \sum_{y\in\pi(o,v)\setminus\{o\}} \left(\log_2 \frac{1}{p(y|y^-)} + \log_2 \frac{1}{q(y|y^-)}\right) \\
&= \sum_{v\in\partial T} \mathbb{P}(v) \left(\log_2 \frac{1}{\mathbb{P}(v)} + \log_2 \frac{1}{\mathbb{Q}(v)}\right) \\
&\le \sum_{v\in\partial T} \mathbb{P}(v) \left(\log_2 \frac{2}{\mathbb{P}(v)} + \left|\log_2 \frac{\mathbb{P}(v)}{\mathbb{Q}(v)}\right|\right) = 2\mathsf{H}(\mathbb{P}) + \sum_{v\in\partial T} \mathbb{P}(v) \left|\log_2 \frac{\mathbb{P}(v)}{\mathbb{Q}(v)}\right|,
\end{aligned}$$

which is finite by assumption. In particular, we get for each internal node x that $|\Delta_{\mathbb{P}}|f(x) < \infty$, that is, $\mathsf{D}_{p,q}(x)$ converges absolutely. □

In the following two lemmas we state some elementary properties of the variational distance between two probability distributions, and the defining function of the entropy.

(3.4) Lemma. *For two probability distributions* $\nu_1\,,\ \nu_2$ *on any finite or countable set* $\mathcal{X}$*, let*

$$\|\nu_1 - \nu_2\|_1 = \sum_{x\in\mathcal{X}} |\nu_1(x) - \nu_2(x)|$$

Then we have the following.

(a) $$\sum_{x\,:\,\nu_1(x)>\nu_2(x)} \big(\nu_1(x)-\nu_2(x)\big) = \sum_{x\,:\,\nu_1(x)<\nu_2(x)} \big(\nu_2(x)-\nu_1(x)\big) = \frac{1}{2}\|\nu_1-\nu_2\|_1\,.$$

(b) [Pinsker's inequality] $$\|\nu_1-\nu_2\|_1^2 \le 2\ln 2\ \mathsf{D}(\nu_1\|\nu_2)\,.$$

The optimal constant $2\ln 2$ is due to Csiszár, Kemperman, and Kullback in independent work of the 1960ies. There are also converse inequalities; see Csiszár and Talata [5], Verdú [11] and Sason [10] and the references given there. The following is elementary.

(3.5) Lemma. *The function $\varphi(t) = t\log_2(1/t)$, $t\in[0\,,\,1]$, with the convention $\varphi(0)=0$, has the following properties.*

(i) *φ is strictly concave and assumes its maximal value $1/(e\ln 2)$ at $t=1/e$.*

(ii) *For $0\le t<t+\delta\le 1$, we have $|\varphi(t+\delta)-\varphi(t)|\le\max\{\varphi(\delta),\varphi(1-\delta)\}$.*

(iii) *For $0\le\delta\le 1/2$, we have $\max\{\varphi(\delta),\varphi(1-\delta)\}=\varphi(\delta)$.*

(iv) *For $t,u\in[0\,,\,1/(2e)]$, we have $\varphi(t+u)\le 2\big(\varphi(t)+\varphi(u)\big)$.*

B. Tightness and uniform summability

Given our tree, suppose that for each internal node x, we have a real valued function $y\mapsto g(y|x)$, defined on the successors y of x. We are interested in the series $\sum_{y\in N(x)} g(y|x)$. Then we say that the family of functions, resp., the associated series are *tight* or *uniformly (absolutely) summable,* if for every $\varepsilon>0$ there is an integer $M_\varepsilon\ge 1$ such that for each $x\in\mathsf{int}\,T$, there is a subset $N_\varepsilon(x)\subset N(x)$ such that

$$|N_\varepsilon(x)|\le M_\varepsilon \quad\text{and}\quad \sum_{y\in N(x)\setminus N_\varepsilon(x)} \big|g(y|x)\big| < \varepsilon\,. \tag{3.6}$$

(The terminology "tight" is commonly used in presence of a family of probability distributions.) As long as all forward degrees as well as the tree are finite, (3.6) is always valid. In this case, we need no ε, i.e., we can replace $N_\varepsilon(x)$ by all of $N(x)$ and M_ε by

$$M=\max\{\mathsf{deg}^+(x) : x\in\mathsf{int}\,T\}. \tag{3.7}$$

When there are only finitely many nodes with infinite forward degree, then (3.6) is equivalent with absolute convergence of the involved series. The same is also true when there are only finitely many types of different functions $g(\cdot|x)$ up to bijections between the sets $N(x)$, $x\in\mathsf{int}\,T$.

In our context, we shall want that all the sums

$$\mathsf{H}_p(x)=\sum_y p(y|x)\log\frac{1}{p(y|x)} \quad\text{and}\quad \mathsf{H}_q(x)=\sum_y q(y|x)\log\frac{1}{q(y|x)} \tag{3.8}$$

are uniformly summable. Uniform summability implies that $\mathsf{H}_q(x)\le\log_2 M_\varepsilon+1/(e\ln 2)+\varepsilon$. In particular, if $\inf\{\ell(x) : x\in\mathsf{int}\,T\}>0$ then $A=\sup\{\mathsf{H}_q(x)/\ell(x) : x\in\mathsf{int}\,T\}<\infty$. This will be needed further below.

We want to formulate sufficient conditions for uniform summability which are easy to grasp. We can enumerate each set $N(x) = \{y_n(x) : n \in \mathbb{N}, n \le \mathsf{deg}^+(x)\}$ (finite or countable) and set

$$p_x(n) = p\big(y_n(x)|x\big) \text{ for } n \le \mathsf{deg}^+(x) \quad \text{and} \quad p_x(n) = 0 \text{ for } n > \mathsf{deg}^+(x)$$

(the latter when $\mathsf{deg}^+(x)$ is finite). Then tightness of the family of distributions $p_x(\cdot)$, $x \in \mathsf{int}\, T$, means that we can do the above enumerations such that

$$\tau(k) = \sup\{\tau_x(k) : x \in \mathsf{int}\, T\} \to 0 \text{ as } k \to \infty\,, \quad \text{where} \quad \tau_x(k) = \sum_{n=k}^{\infty} p_x(n)\,. \tag{3.9}$$

In this case, $p(n) = \tau(n) - \tau(n+1)$ defines a probability distribution on $\mathbb{N}$, which *stochastically dominates* each distribution p_x, that is, its *tails* $\tau(k)$ majorise the tails of all p_x.

(3.10) Lemma. *Suppose that the distributions $p_x = p(\cdot|x)$ satisfy* (3.9) *and that the dominating distribution p has finite mean,*

$$\sum_{n=1}^{\infty} n\, p(n) < \infty\,.$$

Then the entropies $\mathsf{H}_x(p)$, $x \in \mathsf{int}\, T$, are uniformly summable.

If in addition the two probability measures $\mathbb{P}$ and $\mathbb{Q}$ on ∂T are such that

$$\frac{1}{c} \cdot \mathbb{Q} \le \mathbb{P} \le c \cdot \mathbb{Q} \tag{3.11}$$

for some finite $c > 0$, then also the entropies $\mathsf{H}_x(q)$, $x \in \mathsf{int}\, T$, are uniformly summable.

Proof. By Lemma 3.5, the function $\varphi(t)$ is monotone increasing in the interval $[0\,,\, 1/e]$. If $n \ge 2$ and $p_x(n) \le 2^{-n}$ then $\varphi\big(p_x(n)\big) \le n2^{-n}$. Otherwise, $\varphi\big(p_x(n)\big) \le np_x(n)$. Choose $m \ge 2$. We have

$$\sum_{n=m}^{\infty} np_x(n) = m\tau_x(m) + \sum_{k=m+1}^{\infty} \tau_x(k) \le m\tau(m) + \sum_{k=m+1}^{\infty} \tau(k) = \sum_{n=m}^{\infty} np(n)\,.$$

Consequently,

$$\sum_{n=m}^{\infty} \varphi\big(p_x(n)\big) \le \sum_{n=m}^{\infty} n\big(2^{-n} + p(n)\big)\,,$$

which tends to 0 as $m \to \infty$. This proves uniform summability. The second part is obvious. □

4. A comparison theorem

Here is our main result.

(4.1) Theorem. *Let $\mathbb{P}$ and $\mathbb{Q}$ be probability measures supported by all leaves in ∂T, and let $p(\cdot|\cdot)$ and $q(\cdot|\cdot)$ be the transition kernels of the respective associated forward Markov chains. Suppose that $\mathsf{H}(\mathbb{P})$, $\mathsf{H}(\mathbb{Q})$ and $\mathsf{D}(\mathbb{P}\|\mathbb{Q})$ are finite and that the series (sums) $\mathsf{H}_p(x)$ and $\mathsf{H}_q(x)$, $x \in \mathsf{int}\, T$, are uniformly summable (without loss of generality with the same value M_ε and sets $N_\varepsilon(x)$ for every $\varepsilon > 0$).*

Then, for any choice of $\varepsilon > 0$ *and* $0 < \delta < 1/2$,

$$\left|\frac{\mathsf{H}(\mathbb{P})}{\ell(\mathbb{P})} - \frac{\mathsf{H}(\mathbb{Q})}{\ell(\mathbb{Q})}\right| \le L\Big(2\varepsilon + M_\varepsilon\varphi(\delta)\Big) + \frac{C\sqrt{L}}{\delta}\sqrt{\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell(\mathbb{P})}} + \frac{A-a}{2}\,\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_1\,,$$

where

$$L = \frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}\,, \qquad C = \frac{2\sqrt{2}}{e\sqrt{\ln 2}} \approx 1.25\,,$$
$$A = \sup\left\{\frac{\mathsf{H}_q(x)}{\ell(x)} : x \in \mathsf{int}\,T\right\}, \quad \textit{and} \quad a = \inf\left\{\frac{\mathsf{H}_q(x)}{\ell(x)} : x \in \mathsf{int}\,T\right\}.$$

When T *has bounded forward degrees, we can set* $\varepsilon = 0$ *and replace the* M_ε *of* (3.6) *by the* M *of* (3.7).

We postpone a detailed discussion of various aspects and special cases of this inequality to the next section. Right now, we embark on the proof, preceded by a preliminary lemma. We always require the tightness assumptions of Theorem 4.1.

(4.2) Lemma. *For* $\varepsilon > 0$, $0 < \delta \le 1/2$ *and* $x \in \mathsf{int}\,T$,

$$\big|\mathsf{H}_p(x) - \mathsf{H}_q(x)\big| \le 2\varepsilon + M_\varepsilon\varphi(\delta) + \frac{2}{e\ln 2}\frac{1}{\delta}\,\big\|p(\cdot|x) - q(\cdot|x)\big\|_1\,.$$

Proof. If $y \in N(x)$ is such that $\big|p(y|x) - q(y|x)\big| < \delta$, then Lemma 3.5 yields

$$\big|\varphi\big(p(y|x)\big) - \varphi\big(q(y|x)\big)\big| \le \varphi(\delta).$$

In view of this,

$$\big|\mathsf{H}_p(x) - \mathsf{H}_q(x)\big| \le \sum_{y\in N(x)} \big|\varphi\big(p(y|x)\big) - \varphi\big(q(y|x)\big)\big| = \mathrm{Sum}_1 + \mathrm{Sum}_2$$

with

$$\begin{aligned}\mathrm{Sum}_1 &= \sum_{y\,:\,|p(y|x)-q(y|x)|<\delta} \big|\varphi\big(p(y|x)\big) - \varphi\big(q(y|x)\big)\big| \\ &\le \sum_{y\in N(x)\setminus N_\varepsilon(x)} \Big(\big|\varphi\big(p(y|x)\big)\big| + \big|\varphi\big(q(y|x)\big)\big|\Big) + \sum_{y\in N_\varepsilon(x)} \varphi(\delta) \\ &< 2\varepsilon + M_\varepsilon\varphi(\delta)\,.\end{aligned}$$

Next,

$$\begin{aligned}\mathrm{Sum}_2 &= \sum_{y\,:\,|p(y|x)-q(y|x)|\ge\delta} \big|\varphi\big(p(y|x)\big) - \varphi\big(q(y|x)\big)\big| \\ &\le \sum_{y\,:\,|p(y|x)-q(y|x)|\ge\delta} \frac{2\max\varphi}{\delta}\,\delta \\ &\le \frac{2\max\varphi}{\delta} \sum_{y\,:\,|p(y|x)-q(y|x)|\ge\delta} \big|p(y|x) - q(y|x)\big| \\ &\le \frac{2\max\varphi}{\delta}\,\big\|p(\cdot|x) - q(\cdot|x)\big\|_1\,,\end{aligned}$$

completing the proof. □

Proof of Theorem 4.1. We start by observing that

$$\sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}=\frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}. \tag{4.3}$$

Using Lemma 3.1, we write

$$\begin{aligned}\left|\frac{\mathsf{H}(\mathbb{P})}{\ell(\mathbb{P})}-\frac{\mathsf{H}(\mathbb{Q})}{\ell(\mathbb{Q})}\right| &= \left|\sum_{x\in\mathsf{int}\,T}\left(\mu_{\mathbb{P}}(x)\frac{\mathsf{H}_p(x)}{\ell(x)}-\mu_{\mathbb{Q}}(x)\frac{\mathsf{H}_q(x)}{\ell(x)}\right)\right| \\ &\le \sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}\left|\mathsf{H}_p(x)-\mathsf{H}_q(x)\right| \;+\; \left|\sum_{x\in\mathsf{int}\,T}\frac{\mathsf{H}_q(x)}{\ell(x)}\left(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\right)\right| \\ &= \qquad\qquad \mathrm{Sum}_I \qquad\qquad + \qquad\qquad |\mathrm{Sum}_{I\!I}|\end{aligned}$$

For the following sequence of estimates to bound Sum_I, in (1) we use Lemma 4.2, in (2) equation (4.3) and Lemma 3.4(b), in (3) the Cauchy-Schwarz inequality, and in (4) once more (4.3), as well as Proposition 3.3,

$$\begin{aligned}\mathrm{Sum}_I &\overset{(1)}{\le} \sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}\Big(2\varepsilon+M_\varepsilon\,\varphi(\delta)\Big) \;+\; \sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}\,\frac{2\max\varphi}{\delta}\,\big\|p(\cdot|x)-q(\cdot|x)\big\|_1 \\ &\overset{(2)}{=} \frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}\Big(2\varepsilon+M_\varepsilon\varphi(\delta)\Big) \;+\; \frac{C}{\delta}\sum_{x\in\mathsf{int}\,T}\sqrt{\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}}\sqrt{\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}\,\mathsf{D}_{p,q}(x)} \\ &\overset{(3)}{\le} \frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}\Big(2\varepsilon+M_\varepsilon\varphi(\delta)\Big) \;+\; \frac{C}{\delta}\sqrt{\sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}}\sqrt{\sum_{x\in\mathsf{int}\,T}\frac{\mu_{\mathbb{P}}(x)}{\ell(x)}\mathsf{D}_{p,q}(x)} \\ &\overset{(4)}{=} \frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}\Big(2\varepsilon+M_\varepsilon\varphi(\delta)\Big) \;+\; \frac{C}{\delta}\sqrt{\frac{\ell_\sharp(\mathbb{P})}{\ell(\mathbb{P})}}\sqrt{\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell(\mathbb{P})}}.\end{aligned}$$

Next, by Lemma 3.4(a),

$$\begin{aligned}&\sum_{x\in\mathsf{int}\,T}\frac{\mathsf{H}_q(x)}{\ell(x)}\left(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\right) \\ &\qquad\le \sum_{x\,:\,\mu_{\mathbb{P}}(x)>\mu_{\mathbb{Q}}(x)}A\left(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\right)+\sum_{x\,:\,\mu_{\mathbb{P}}(x)<\mu_{\mathbb{Q}}(x)}a\left(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\right) \\ &\qquad= \frac{A-a}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1\,.\end{aligned}$$

Exchanging the roles of $\mu_{\mathbb{P}}$ and $\mu_{\mathbb{Q}}$, we get the same upper bound. This proves that $|\mathrm{Sum}_{I\!I}|\le\frac{A-a}{2}\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1$, concluding the proof. □

5. Discussion and variations

Our way of thinking is that $\mathbb{Q}$ is the basic reference measure, while $\mathbb{P}$ may vary; it should be considered as a perturbation of $\mathbb{Q}$.

We say that two internal nodes x and x' have the *same* $\mathbb{Q}$*-type,* if there is a bijection $\pi : N(x)\to N(x')$ such that

$$q(\pi y|x')=q(y|x)\quad\text{for all}\;\; y\in N(x)\,. \tag{5.1}$$

A. Identical $\mathbb{Q}$-types

The simplest, but most significant situation is the one where all internal nodes have the same $\mathbb{Q}$-type. In this case, $\mathsf{H}_q(x) = \mathsf{H}_q$ is constant, and uniform summability of $\mathsf{H}_q(x)$ reduces to finiteness of H_q. If in addition $\ell = \ell_\sharp$, then $A = a$, and by Lemma 3.1, one has $\mathsf{H}(\mathbb{Q}) = \ell_\sharp(\mathbb{Q}) \cdot \mathsf{H}_q$. The inequality of Theorem 4.1 becomes

$$\left| \frac{\mathsf{H}(\mathbb{P})}{\ell_\sharp(\mathbb{P})} - \mathsf{H}_q \right| \le 2\varepsilon + M_\varepsilon \varphi(\delta) + \frac{C}{\delta} \sqrt{\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell_\sharp(\mathbb{P})}}\,. \tag{5.2}$$

When T is finite, so that we can set $\varepsilon = 0$ and replace M_ε by M, this is the main result of [1]. (5.2) provides an extension to the case when all $N(x)$ are infinite and there is one $\mathbb{Q}$-type with $\mathsf{H}_q < \infty$. Note that when in addition the measures $\mathbb{P}$ and $\mathbb{Q}$ are comparable as in (3.11) then the uniform summability conditions of Theorem 4.1 hold.

Suppose furthermore that T has height n, that is, $|v|_\sharp = n$ for each leaf of T. Then with suitable labelling, $\mathbb{Q} = q^{\otimes n} = q \otimes \cdots \otimes q$ is the n-fold product measure of $q(x) := q(x|o)$, $x \in N(o)$, or in other words, it is the joint distribution of n independent random variables $X_1, \dots, X_n$ with common distribution $q(\cdot)$, and (5.2) becomes

$$\left| \frac{\mathsf{H}(\mathbb{P})}{n} - \mathsf{H}_q \right| \le 2\varepsilon + M_\varepsilon \varphi(\delta) + \frac{C}{\delta} \sqrt{\frac{\mathsf{D}(\mathbb{P}\|q^{\otimes n})}{n}}\,. \tag{5.3}$$

Here we should think of $\mathbb{P} = \mathbb{P}^{(k)}$ as the joint distribution of a perturbation of $(X_1, \dots, X_n)$ such that for fixed n, we have that $\mathsf{D}(\mathbb{P}^{(k)}\|q^{\otimes n}) \to 0$ as $k \to \infty$. Then, under the stated conditions of uniform summabilty – in particular, when $\frac{1}{c} \cdot q^{\otimes n} \le \mathbb{P}^{(k)} \le c \cdot q^{\otimes n}$ –

$$\lim_{k\to\infty} \frac{\mathsf{H}(\mathbb{P}^{(k)})}{n} = \mathsf{H}_q\,.$$

We shall come back to this further below, considering $n \to \infty$.

To put (5.3) into perspective, consider the simple case where also $\mathbb{P} = p^{\otimes n}$. Then

$$\frac{\mathsf{H}(\mathbb{P})}{n} - \frac{\mathsf{H}(\mathbb{Q})}{n} = \mathsf{H}\big(p(\cdot)\big) - \mathsf{H}\big(q(\cdot)\big) \quad \text{and} \quad \frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{n} = \mathsf{D}\big(p(\cdot)\|q(\cdot)\big)\,,$$

and when q happens to be uniformly distributed (on a finite set) then the two quantities coincide.

B. Variable $\mathbb{Q}$-types

If $\mathsf{H}_q(x)$ is not constant, then we may use the length function $\ell_q(x) = \mathsf{H}_q(x)$ instead of $\ell_\sharp$. Again, $A = a$, and (3.2) leads to

$$\left| \frac{\mathsf{H}(\mathbb{P})}{\ell_q(\mathbb{P})} - 1 \right| \le 2\varepsilon + M_\varepsilon \varphi(\delta) + \frac{C}{\delta} \sqrt{\frac{\mathsf{D}(\mathbb{P}\|\mathbb{Q})}{\ell_q(\mathbb{P})}}\,. \tag{5.4}$$

If $\mathsf{H}_q(x)$ varies between two positive bounds, then $\ell_q(\mathbb{P})/\ell_\sharp(\mathbb{P})$ varies between the same bounds, i.e., the two expected lenghts are of the same order of magnitude. This applies, in particular, to the situation when there are only finitely many different $\mathbb{Q}$-types: this is the case, for example, when the tree encodes the trajectories of a finite state Markov chain; compare with (6.7.b) below.

Otherwise, e.g. if we want to stick to $\ell_\sharp$, we need to discuss the quantity $\frac{A-a}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1$, since it remains present in the estimate of Theorem 4.1. First of all, we have the following upper bound:

$$\frac{A-a}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1 \le \frac{A-a}{\ell(\mathbb{P})}\sum_{v\in\partial T}|v|_\ell\,\big|\mathbb{P}(v)-\mathbb{Q}(v)\big|\,. \tag{5.5}$$

Second, we can use the following variant.

(5.6) Lemma. *Let S be a cross section of T (Def. 2.3), and let*

$$A^*_S=\sup\left\{\frac{\mathsf{H}_q(x)}{\ell(x)}:x\in\mathsf{int}\,T\setminus\mathsf{int}\,T^S\right\},\quad \textit{and}\quad a^*_S=\inf\left\{\frac{\mathsf{H}_q(x)}{\ell(x)}:x\in\mathsf{int}\,T\setminus\mathsf{int}\,T^S\right\}.$$

Then the term $\frac{A-a}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1$ in the estimate of Theorem 4.1 can be replaced with

$$\frac{A^*_S-a^*_S}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1+(A-a)\,\max\left\{\frac{\ell(\mathbb{P}_S)}{\ell(\mathbb{P})}\,,\,\frac{\ell(\mathbb{Q}_S)}{\ell(\mathbb{Q})}\right\}.$$

This is useful, if $H_q(x)/\ell(x)$ is "almost constant" outside a small subtree. By this we mean that there is a cross section S such that $A^*_S-a^*_S$ is small, and at the same time, T is much larger that T^S, so that also $\ell(\mathbb{P}_S)/\ell(\mathbb{P})$ and $\ell(\mathbb{Q}_S)/\ell(\mathbb{Q})$ are small. See below in §6.

Proof of Lemma 5.6. We need to modify the estimate of the term $|\mathrm{Sum}_{I\!I}|$ in the proof of Theorem 4.1. Consider the sets $\mathsf{Pos}=\{x\in\mathsf{int}\,T:\mu_{\mathbb{P}}(x)>\mu_{\mathbb{Q}}(x)\}$ and $\mathsf{Neg}=\{x\in\mathsf{int}\,T:\mu_{\mathbb{P}}(x)>\mu_{\mathbb{Q}}(x)\}$, as well as $\mathsf{Pos}_S=\mathsf{Pos}\cap\mathsf{int}\,T^S$ and $\mathsf{Neg}_S=\mathsf{Neg}\cap\mathsf{int}\,T^S$. We split

$$\begin{aligned}
&\sum_{x\in\mathsf{int}\,T}\frac{\mathsf{H}_q(x)}{\ell(x)}\,\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)\\
&\quad\le\sum_{x\in\mathsf{Pos}_S}A\,\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)+\sum_{x\in\mathsf{Neg}_S}a\,\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)\\
&\qquad+\sum_{x\in\mathsf{Pos}\setminus\mathsf{Pos}_S}A^*_S\,\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)+\sum_{x\in\mathsf{Neg}\setminus\mathsf{Neg}_S}a^*_S\,\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)\\
&\quad=A^*_S\sum_{x\in\mathsf{Pos}}\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)+a^*_S\sum_{x\in\mathsf{Neg}}\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)\\
&\qquad+(A-A^*_S)\sum_{x\in\mathsf{Pos}_S}\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)+(a-a^*_S)\sum_{x\in\mathsf{Neg}_S}\big(\mu_{\mathbb{P}}(x)-\mu_{\mathbb{Q}}(x)\big)\\
&\quad\le\frac{A^*_S-a^*_S}{2}\,\|\mu_{\mathbb{P}}-\mu_{\mathbb{Q}}\|_1+(A-A^*_S)\,\mu_{\mathbb{P}}(\mathsf{Pos}_S)+(a^*_S-a)\mu_{\mathbb{Q}}(\mathsf{Neg}_S)\,.
\end{aligned}$$

Now, by the definition (2.4) of $\mu_{\mathbb{P}}$,

$$\mu_{\mathbb{P}}(\mathsf{Pos}_S)=\frac{\ell(\mathbb{P}_S)}{\ell(\mathbb{P})}\,\mu_{\mathbb{P}_S}(\mathsf{Pos}_S)\le\frac{\ell(\mathbb{P}_S)}{\ell(\mathbb{P})},\quad\text{and analogously}\quad\mu_{\mathbb{Q}}(\mathsf{Neg}_S)\le\frac{\ell(\mathbb{Q}_S)}{\ell(\mathbb{Q})}\,.$$

This and an exchange of the roles of $\mathbb{P}$ and $\mathbb{Q}$ leads to the stated upper bound. □

At this point we also remark that one can of course refrain from deriving an upper bound on Sum_I that involves Pinsker's inequality, as in the proof of Theorem 4.1: instead one can work directly with Sum_I. However, recall from above that there are converses

to Pinsker's inequality. We conclude this section with an example which shows that the term $\frac{A-a}{2}\,\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_1$ is indispensable in the estimate of Theorem 4.1.

(5.7) Example. We consider a tree $T = T^{(n)}$ of height $n+1$ whose root has 2 forward neighbours, so that the two branches T_1 and T_2 of height n which are attached to o have all inner nodes with forward degrees d_1 in T_1 and d_2 in T_2, respectively. Then $\partial T = \partial T_1 \cup \partial T_2$. For $0 < \theta < 1$, we define

$$\mathbb{P}_\theta(v) = \mathbb{P}_\theta^{(n)}(v) = \begin{cases} \theta/d_1^n\,, & \text{if } v \in \partial T_1\,, \\ (1-\theta)/d_2^n\,, & \text{if } v \in \partial T_2\,. \end{cases}$$

With $\ell = \ell_\sharp$, we have of course $\ell(\mathbb{P}_\theta) = n+1$. It is easy to compute

$$\mathsf{H}(\mathbb{P}_\theta) = \mathsf{H}(\theta, 1-\theta) + \theta\,\log_2 d_1 + (1-\theta)\,\log_2 d_2\,,$$

where $\mathsf{H}(\theta, 1-\theta)$ is the entropy of the Bernoulli distribution with parameter θ.

Now we take $\mathbb{P} = \mathbb{P}_\theta$, where $\theta \neq 1/2$, and $\mathbb{Q} = \mathbb{P}_{1/2}$. Then it is also easy to compute

$$\mathsf{D}(\mathbb{P}\|\mathbb{Q}) = 1 - H(\theta, 1-\theta)$$

As $n \to \infty$, we have that

$$\frac{\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)}{n} \to 0\,, \quad \text{while} \quad \left|\frac{\mathsf{H}(\mathbb{P}_n)}{n} - \frac{\mathsf{H}(\mathbb{Q}_n)}{n}\right| \to \left|\left(\theta - \tfrac{1}{2}\right)\log_2 d_1 + \left(1 - \theta - \tfrac{1}{2}\right)\log_2 d_2\right|,$$

and the latter limit is non-zero for appropriate choices of θ, d_1 and d_2. Thus, there can be no improved general estimate that can remove the term $\frac{A-a}{2}\,\|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_1$ or replace it by something of significantly smaller order of magnitude. □

The last example can be seen in terms of two stationary stochastic processes on denumerable state spaces. We first toss a coin to decide which of the two processes to run, and then "forget" about the other one. The reference measure $\mathbb{Q}$ corresponds to a fair coin toss, while $\mathbb{P}$ corresponds to a biased one. Then we compare their entropy rates in terms of their Kullback-Leibler divergence. Entropy rates will be featured in the next section.

6. Trees with infinite geodesics

We now drop the assumption that our tree has no infinite geodesics. In their presence, let us describe the boundary ∂T of T. It consists of eventual leaves (vertices $v \neq o$ with no forward neighbour) plus the *ends* of T. Each end w is represented by an infinite geodesic $\pi(o, w)$ starting at o (but when we have antoher geodesic which differs from the given one only by a finite initial piece, then it represents the same end). We can put a metric on $\widehat{T} = \mathsf{int}\, T \cup \partial T$: for any pair of distinct points $v, w \in \widehat{T}$, their *confluent* $v \wedge w$ is the last common vertex on the geodesics $\pi(o, v)$ and $\pi(o, w)$. Then the metric is

$$d(v, w) = \begin{cases} 0\,, & \text{if } v = w\,, \\ 2^{-|v \wedge w|}\,, & \text{if } v \neq w\,. \end{cases}$$

Thus, a sequence (x_n) of vertices of T converges to an end w, if $|x_n \wedge w| \to \infty$. For more details in a context close to the present note, see e.g. [12, Ch.9.B]. Note that in many typical examples, such as the infinite binary tree, ∂T is a Cantor set, whence uncountable.

For $x \in \mathsf{int}\, T$, the branch T_x and its boundary ∂T_x are defined as before. The boundary now may contain ends as well as leaves, and the sets of all ∂T_x , $x \in \mathsf{int}\, T$, are the open (and also closed!) balls for the metric on ∂T. Any probability measure $\mathbb{P}$ on ∂T is given via consistent definition of the values $\mathbb{P}(\partial T_x)$, $x \in \mathsf{int}\, T$. That is, we need to have

$$\mathbb{P}(\partial T) = 1 \quad \text{and} \quad \sum_{y \in N(x)} \mathbb{P}(\partial T_y) = \mathbb{P}(\partial T_x) \ \text{ for every } x \in \mathsf{int}\, T\,.$$

(6.1) Example. The most typical class of examples arises from a stochastic process $(X_n)_{n\ge1}$ on a finite or countable state space $\mathfrak{S}$. In this case, for $n \in \mathbb{N}$, let

$$p_n(s_1, \dots, s_n) = \mathbb{P}[X_1 = s_1\,, \dots, X_n = s_n]\,, \quad s_1, \dots, s_n \in \mathfrak{S}$$

be the joint distribution of $(X_1, \dots, X_n)$. Then we let (the vertex set of) our tree be

$$T = \{s_1 s_2 \cdots s_n : n \ge 0\,,\ s_i \in \mathfrak{S}\,,\ p_n(s_1, \dots, s_n) > 0\}\,. \tag{6.2}$$

For $n = 0$, this includes the *empty word* o, for which $p_0(o) = 1$. The predecessor vertex of $x = s_1 \cdots s_n$ $(n \ge 1)$ is $x^- = s_1 \cdots s_{n-1}$, and

$$p(x|x^-) = \mathbb{P}[X_1 = s_1\,, \dots, X_{n-1} = s_{n-1}\,, X_n = s_n \mid X_1 = s_1\,, \dots, X_{n-1} = s_{n-1}]\,,$$

while for $x = s \in N(o)$, we have $p(x|o) = \mathbb{P}[X_1 = x]$. In this case, T has no leaves, the end space

$$\partial T = \{s_1 s_2 s_3 \cdots : p_n(s_1, \dots, s_n) > 0 \text{ for each } n\}$$

is the space of (active) trajectories of our process, and $\mathbb{P}$ is the distribution of that process. □

In general, we cannot speak directly about the entropy of $\mathbb{P}$ on ∂T. We can take any section S of T and consider $\mathsf{H}(\mathbb{P}_S)$, as in Definition 2.3 and the subsequent lines. We restrict attention to the following sequence of sections with respect to the standard length function:

$$S(n) = \{x \in T : |x| = n\} \cup \{v \in \partial T : |v| < n\}.$$

The second part appears when T has leaves. We write $\mathbb{P}_n = \mathbb{P}_{S(n)}$ and $\mathbb{Q}_n = \mathbb{Q}_{S(n)}$ and assume that $\mathsf{H}(\mathbb{P}_n)$ and $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)$ are finite (absolutely convergent) for each n. Then we can speak about the quantity

$$h = h\big(\mathbb{P}, \ell\big) = \lim_{n\to\infty} \frac{\mathsf{H}(\mathbb{P}_n)}{\ell(\mathbb{P}_n)}\,,$$

whenever the limit exists. This is a variant of entropy rate (asymptotic entropy). In particular, in the situation of Example 6.1, for the standard length function, $h(\mathbb{P}) = h(\mathbb{P}, \sharp)$ is indeed the classical entropy rate of the process (X_n) with distribution $\mathbb{P}$. In any case, relative entropy is well defined; see Gray [6, §7.1]:

$$\mathsf{D}(\mathbb{P}\|\mathbb{Q}) = \lim_{n\to\infty} \mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n) \tag{6.3}$$

Indeed, one can apply the well-known log-sum inequality – see Cover and Thomas [4, Thm. 2.7.1] – to verify that $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)$ is increasing with n. Of course, the limit may be infinite.

(6.4) Theorem. *Let T be a countable tree with infinite geodesics and $\mathbb{P}$ and $\mathbb{Q}$ two probability measures giving positive mass to the set of ends of T and satisfying our tightness requirements for the local entropies at the inner vertices. Assume that*

(i) *The length function ℓ is comparable with $\ell_\sharp$, that is $0 < c_1 \le \ell(x) \le c_2 < \infty$ for suitable c_1, c_2 and all $x \in \mathsf{int}\, T$,*

(ii) *the limit $h = \lim_{|x|\to\infty} \mathsf{H}_q(x)/\ell(x)$ exists and is finite, and*

(iii) *the relative entropy $\mathsf{D}(\mathbb{P}\|\mathbb{Q})$ is finite.*

Then

$$h(\mathbb{P}, \ell) = h(\mathbb{Q}, \ell) = h\,.$$

Proof. Write $\partial_\infty T$ for the space of ends of T. By assumption, $\mathbb{P}(\partial_\infty T) > 0$. Therefore

$$\ell_\sharp(\mathbb{P}_n) \ge n\,\mathbb{P}(\partial_\infty T) \to \infty\,, \quad \text{as } n \to \infty\,. \tag{6.5}$$

The same applies to $\mathbb{Q}$.

Next, we use Lemma 5.6. We start with an arbitrary $\varepsilon > 0$. By (ii) there is an index k such that $A^*_{S(k)} - a^*_{S(k)} < \varepsilon$. We take $n > k$ and consider the subtree $T^{S(n)}$ with $\mathbb{P}_n$ and $\mathbb{Q}_n$. We note that $L = \ell_\sharp(\mathbb{P}_n)/\ell(\mathbb{P}_n) \le c_2$ by assumption (i). Next, we set $\delta_n = \ell(\mathbb{P}_n)^{-1/4}$, which tends to 0 as $n \to \infty$ again by assumption (i), together with (6.5). Now the upper bound of Theorem 4.1, combined with the variant provided by Lemma 5.6, yields

$$\begin{aligned}\left|\frac{\mathsf{H}(\mathbb{P}_n)}{\ell(\mathbb{P}_n)} - \frac{\mathsf{H}(\mathbb{Q}_n)}{\ell(\mathbb{Q}_n)}\right| &< c_2\Big(2\varepsilon + M_\varepsilon\varphi(\delta_n)\Big) \;+\; C\sqrt{c_2}\,\mathsf{D}(\mathbb{P}\|\mathbb{Q})\,\delta_n \\ &\qquad + \varepsilon \;+\; (A-a)\,\max\left\{\frac{\ell(\mathbb{P}_k)}{\ell(\mathbb{P}_n)}\,,\,\frac{\ell(\mathbb{Q}_k)}{\ell(\mathbb{Q}_n)}\right\}\,.\end{aligned}$$

If we let $n \to \infty$, again using assumption (i) together with (6.5) to control the last maximum, we see that

$$\limsup_{n\to\infty}\left|\frac{\mathsf{H}(\mathbb{P}_n)}{\ell(\mathbb{P}_n)} - \frac{\mathsf{H}(\mathbb{Q}_n)}{\ell(\mathbb{Q}_n)}\right| \le (2c_2+1)\,\varepsilon\,,$$

because all other terms tend to 0. We are left with showing that $\mathsf{H}(\mathbb{Q}_n)/\ell(\mathbb{Q}_n) \to h$. This is standard: with ε and k as above,

$$\begin{aligned}\left|\frac{\mathsf{H}(\mathbb{Q}_n)}{\ell(\mathbb{Q}_n)} - h\right| &\le \sum_{x\in\mathsf{int}\,T^{S(n)}}\left|\frac{\mathsf{H}_q(x)}{\ell(x)} - h\right|\,\mu_{\mathbb{Q}_n}(x) \\ &< \sum_{x\in\mathsf{int}\,T^{S(k)}}(A+h)\,\mu_{\mathbb{Q}_n}(x) + \sum_{x\in\mathsf{int}\,T^{S(n)}\setminus T^{(Sk)}}\varepsilon\,\mu_{\mathbb{Q}_n}(x) \le (A+h)\frac{\ell(\mathbb{Q}_k)}{\ell(\mathbb{Q}_n)} + \varepsilon.\end{aligned}$$

This becomes smaller than 2ε when n is large enough. □

(6.6) Remark. In Theorem 6.4, instead of finiteness of $\mathsf{D}(\mathbb{P}\|\mathbb{Q})$ it suffices to assume that $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)/n \to 0$ as $n \to \infty$.

(6.7) Application. We take up Example 6.1, with the difference that we now write $\mathbb{Q}$ for the distribution on the trajectory space of our random process (X_n). We have identified that space with the space of ends of a rooted tree T which has no leaves. We take the natural length function $\ell_\sharp$ and assume that $H_q(x) = h$ is constant. Then the standard

entropy rate $h(\mathbb{Q})$ exists and coincides with h. Now we consider a perturbation $(\widetilde{X}_n)$ of (X_n) with overall distribution $\mathbb{P}$. In this situation, Theorem 6.4 yields that if $\mathsf{D}(\mathbb{P}\|\mathbb{Q}) < \infty$ then also the entropy rate of $(\widetilde{X}_n)$ exists and is equal to the same h. Typical cases where this can be used are as follows.

(a) As in (5.3), let (X_n) be a sequence of i.i.d. discrete random variables with common distribution $q(\cdot)$. Then $\mathbb{Q} = q^{\otimes\mathbb{N}}$ is the infinite product measure of countably many copies of q, and the entropy rate is $h(\mathbb{Q}) = \mathsf{H}(q)$. If we have a perturbed process $(\widetilde{X}_n)$ whose overall distribution $\mathbb{P}$ satisfies $\mathsf{D}(\mathbb{P}_n\|q^{\otimes n})/n \to 0$ then also $h(\mathbb{P}) = \mathsf{H}(q)$.

The following specific example is related to the famous paper of Kakutani [7] on absolute continuity versus orthogonality of infinite product measures. We let $q(\cdot)$ be uniform distribution on $\{1, \dots, M\}$ and $\mathbb{Q} = q^{\otimes\mathbb{N}}$. The associated tree is the M-ary tree, where each vertex has M successors. For any $\alpha \in (0\,,\,1)$, we define $p_\alpha(\cdot)$ on $\{1, \dots, M\}$ by

$$p_\alpha(1) = \frac{1+\alpha}{M}\,,\ p_\alpha(2) = \frac{1-\alpha}{M}\,, \quad \text{and} \quad p_\alpha(s) = \frac{1}{M} \ \text{ for } s = 3, \dots, M\,.$$

Then

$$\mathsf{D}(p_\alpha\|q) = \mathsf{H}(q) - \mathsf{H}(p_\alpha) = \frac{f(\alpha)}{M}\,, \quad \text{where} \quad f(\alpha) = (1+\alpha)\,\log_2(1+\alpha) + (1-\alpha)\,\log_2(1-\alpha).$$

We note that $f(\alpha)/\alpha^2 \to \ln 2$ as $\alpha \to 0$. Now we take a sequence (α_n) and the perturbed measure $\mathbb{P} = \bigotimes_{n=1}^{\infty} p_{\alpha_n}$. Then

$$\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n) = \mathsf{H}(\mathbb{Q}_n) - \mathsf{H}(\mathbb{P}_n) = \frac{f(\alpha_1) + \cdots + f(\alpha_n)}{M}$$

If $\alpha_n \to 0$ then $\frac{1}{n}\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n) \to 0$, whence $h(\mathbb{P}) = h(\mathbb{Q}) = \log_2 M$. If in particular $\alpha_n = n^{-\beta}$ with $\beta > 0$ then we see that $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n) = \infty$ precisely when $\beta < 1/2$.

(b) Now suppose that (X_n) is a Markov chain on a finite or countable state space $\mathfrak{S}$ with transition matrix $\mathsf{Q} = \big(\mathsf{q}(s'|s)\big)_{s,s'\in\mathfrak{S}}$ and initial distribution ν on $\mathfrak{S}$. Denote once more by $\mathbb{Q}$ – indstead of $\mathbb{P}$ as in Example 6.1 – the overall probability distribution of the process on its trajectory space, which we interpret as the boundary at infinity of our tree T. In this case, as in (6.1) but with q_n instead of p_n, we have

$$q_n(s_1\,, \dots, s_n) = \mathsf{q}(s_n|s_{n-1}) \cdots \mathsf{q}(s_2|s_1)\nu(s_1)\,.$$

In particular, for a node $x = s_1 \cdots s_{n-1}s_n$ ouf our tree, with $n \geq 1$, we have $q(x|x^-) = \mathsf{q}(s_n|s_{n-1})$. Now suppose that the tranisition matrix is such that all its rows $\mathsf{q}(\cdot|s)$ have the same finite entropy h. Then $h(\mathbb{Q}) = h$. Again, if we have a perturbation $(\widetilde{X}_n)$ of (X_n) with overall distribution $\mathbb{P}$ such that our tightness requirements are met and $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)/n \to 0$, then also $h(\mathbb{P}) = h$.

(c) A specific example concerning the last situation is that of a finite or infinite oriented graph with vertex set $\mathfrak{S}$ which is d-regular, that is, every vertex has d outgoing edges. For our Markov chain, we take *simple random walk* which at each vertex moves with equal probability to one of those neighbours. Then $h = \log_2 \mathsf{d}$ and all of the above applies.

There are many further examples of Markov chains where $\mathsf{H}\big(\mathsf{q}(\cdot|s)\big)$ is constant, in particular those where all rows $\mathsf{q}(\cdot|s)$ of the transition matrix are permutations of each

other. If more generally $H_q(x)$ is not constant, then we can use $\ell(x) = H_q(x)$ as in (5.4), and when $\mathsf{D}(\mathbb{P}\|\mathbb{Q}) < \infty$ we get that

$$a \le \liminf_{n\to\infty} \frac{1}{n} H(\widetilde{X}_1, \dots, \widetilde{X}_n) \le \limsup_{n\to\infty} \frac{1}{n} H(\widetilde{X}_1, \dots, \widetilde{X}_n) \le A\,,$$

where a and A are as in Theorem 4.1. In this situation, it may also be of interest to note that if $h(\mathbb{P}) = \lim_{n\to\infty} \frac{1}{n} H(\widetilde{X}_1, \dots, \widetilde{X}_n)$ exists, then

$$h(\mathbb{P}) = \lim_{n\to\infty} \frac{\ell_q(\mathbb{P}_n)}{n}\,,$$

which follows once more from (5.4). □

(6.8) Remark. All results stated so far adapt easily to the situation where the reference measure $\mathbb{Q}$ is supported by the whole of ∂T, while the probability measure $\mathbb{P}$ is supported by a possibly strict subset. In this situation, we may have $p(x|x^-) = 0$ for some edges of T, and the node-average measure $\mu_{\mathbb{P}}$ will only live on the induced subtree of T which is spanned by all vertices that can be reached from the root.

The following is our final application.

Random perturbations. We start once more with the setting of Example 6.1, that is, the infinite tree is as in (6.2), and its boundary is equipped with our reference probability measure $\mathbb{Q}$. At the internal vertices we have the transition probabilities $q(\cdot|x)$ and a length function $\ell(x)$.

$\mathbb{Q}$ describes the basic stochastic process which hovever is randomly perturbed: there is a second collection of transition probabilities $q'(\cdot|x)$ at each internal vertex, supported by some or all forward neighbours of x. We assume that both families of probability distributions $q(\cdot|x)$ and $q'(\cdot|x)$, $x \in T$, are tight and that their entropies are uniformly summable; see Lemma 3.10 for a sufficient condition.

Now we consider a sequence of *random variables* $\boldsymbol{\delta}_n$, $n \ge 0$, taking their values in the interval $[0\,,\,1]$, defined on a separate probablity space. Our perturbed random process is given via the forward transition probabilities

$$p(y|x) = (1 - \boldsymbol{\delta}_n)\, q(y|x) + \boldsymbol{\delta}_n\, q'(y|x)\,, \quad \text{if } \ y \in S(n)\,. \tag{6.9}$$

Note that this defines *random* transition probabilities, so that we get a randomized process: at time n – when $x \in S(n)$ – the moving particle chooses to proceed according to the "wrong" forward transition probabilites $q'(\cdot|x)$ with probability $\boldsymbol{\delta}_n$, while it follows the "correct" transition rule with the complementary probability. We write $\mathbb{P}$ for the *random* probability measure on ∂T induced by $p(\cdot|\cdot)$, and as before $\mathbb{P}_n = \mathbb{P}_{S(n)}$.

(6.10) Theorem. *In the setting of* (6.9) *with tight families of forward transition probabilities $q(\cdot|\cdot)$ and $q'(\cdot|\cdot)$ as well as uniformly summable local entropies, assume that*

(i) *The length function ℓ is comparable with $\ell_\sharp$,*

(ii) *the limit $h = \lim_{|x|\to\infty} \mathsf{H}_q(x)/\ell(x)$ exists and is finite, and*

(iii) *there is $K < \infty$ such that $\mathsf{D}_{q',q}(x) \le K$ for all $x \in T$.*

Under these conditions, if $\lim_{n\to\infty} \mathbb{E}(\boldsymbol{\delta}_n) = 0$ *then* $\lim_{n\to\infty} \dfrac{\mathsf{H}(\mathbb{P}_n)}{\ell(\mathbb{P}_n)} = h$ *in probability.*

Moreover, if $\boldsymbol{\delta}_n \to 0$ *almost surely, then* $h(\mathbb{P}, \ell) = h$ *almost surely, i.e., for almost every realisation of the sequence* $(\boldsymbol{\delta}_n)$*.*

Proof. First of all, we need to verify that the (random) local entropies $\mathsf{H}_p(x)$ are uniformly summable. By tightness of $q(\cdot|x)$ and $q'(\cdot|x)$, there is $\widetilde{M}$ such that for each x there is a set $\widetilde{N}(x) \subset N(x)$ with $|\widetilde{N}(x)| \le \widetilde{M}$ such that

$$\max\{q(y|x), q'(y|x)\} \le 1/(2e) \quad \text{for all} \quad y \in N(x) \setminus \widetilde{N}(x)\,.$$

Using Lemma 3.5(iv),

$$\sum_{y\in N(x)\setminus\widetilde{N}(x)} \varphi\big(p(y|x)\big) \le 2 \sum_{y\in N(x)\setminus\widetilde{N}(x)} \Big(\varphi\big(q(y|x)\big) + \varphi\big(q'(y|x)\big)\Big)$$

As x varies, the right hand sides are uniformly summable along with $\mathsf{H}_q(x)$ and $\mathsf{H}_{q'}(x)$.

At this point, with the same line of reasoning as in the proof of Theorem 6.4, what remains is to show that $\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)/\ell(\mathbb{P}_n) \to 0$ in probability, resp. almost surely. For this purpose, we first take $x \in S(n)$, and use convexity, see [4, Thm. 2.7.2]:

$$\begin{aligned}\mathsf{D}_{p,q}(x) &= \mathsf{D}\Big((1-\boldsymbol{\delta}_n)\,q(\cdot|x) + \boldsymbol{\delta}_n\,q'(\cdot|x)\Big\|(1-\boldsymbol{\delta}_n)\,q(\cdot|x) + \boldsymbol{\delta}_n\,q(\cdot|x)\Big)\\ &\le \boldsymbol{\delta}_n\,\mathsf{D}\big(q'(\cdot|x)\big\|q(\cdot|x)\big) \le \boldsymbol{\delta}_n\,K\,.\end{aligned}$$

Now we combine the last inequality with Proposition 3.3 and the Definition (2.4) of $\mu_{\mathbb{P}_n}$,

$$\frac{\mathsf{D}(\mathbb{P}_n\|\mathbb{Q}_n)}{\ell(\mathbb{P}_n)} \le \mathbb{E}_{\mu_{\mathbb{P}_n}}\left(\frac{\boldsymbol{\delta}_{|x|}\,K}{\ell(x)}\right) = \sum_{k=0}^{n}\sum_{x\in S(k)} \mu_{\mathbb{P}_n}(x)\frac{\boldsymbol{\delta}_k\,K}{\ell(x)} = \frac{D}{\ell(\mathbb{P}_n)}\sum_{k=0}^{n}\boldsymbol{\delta}_k \le \frac{K}{c_1 n}\sum_{k=0}^{n}\boldsymbol{\delta}_k\,.$$

The last inequality follows from assumption (i). If $\mathbb{E}(\boldsymbol{\delta}_n) \to 0$, resp. $\boldsymbol{\delta}_n \to 0$ almost surely then

$$\frac{1}{n}\sum_{k=0}^{n}\boldsymbol{\delta}_k \to 0$$

in probability, resp. almost surely. □

(6.11) Example. Suppose that $\boldsymbol{\delta}_n$ is a Bernoulli random variable with $\mathbb{P}[\boldsymbol{\delta}_n = 1] = 1/n^\beta$ and $\mathbb{P}[\boldsymbol{\delta}_n = 0] = 1 - 1/n^\beta$, where $\beta > 0$. Then $\mathbb{E}(\boldsymbol{\delta}_n) \to 0$, so that we get convergence in probability in Theorem 6.10. If $\beta > 1$ then by the Borel-Cantelli Lemma we have with probability 1 that $\boldsymbol{\delta}_n = 0$ for all but a (random) finite number of n. In this case, we get that $h(\mathbb{P}, \ell) = h(\mathbb{Q}, \ell)$ almost surely.

References

[1] Böcherer, G,. and Amjad, R. A.: *Informational divergence and entropy rate on rooted trees with probabilities,* IEEE Int. Symp. Inf. Theory, June 2014; `arXiv:1310.2882`

[2] Ciuperca G., and Girardin V.: *On the estimation of the entropy rate of finite Markov chains,* Proceedings XIth International Symposium on ASMDA May 2005, Brest, France, pp. 1109–1117; `https://conferences.telecom-bretagne.eu/asmda2005/`

[3] Ciuperca, G., Girardin, V., and Lhote, L.: *Computation and estimation of generalized entropy rates for denumerable Markov chains,* IEEE Trans. Inform. Theory **57** (2011) 4026–4034.

[4] Cover, T. M., and Thomas, J. A.: *Elements of Information Theory,* 2nd ed., Wiley, Hoboken, 2006.

[5] Csiszár, I., and Talata, Z.: *Context tree estimation for not necessarily finite memory processes, via BIC and MDL,* IEEE Trans. Inform. Theory **52** (2006) 1007–1016.
[6] Gray, R. M.: *Entropy and information theory,* 2nd ed., Springer, New York, 2011.
[7] Kakutani, S.: *On equivalence of infinite product measures,* Ann. of Math. (2) **49** (1948) 214–224.
[8] Massey, J. L.: *Applied Digital Information Theory I,* lecture notes, ETH Zürich, online: `http://www.isiweb.ee.ethz.ch/archive/massey_scr/adit1.pdf`
[9] Rueppel, R. A., and Massey, J. L.: *Leaf-average node-sum interchanges in rooted trees with applications,* in *Communications and Cryptography: two Sides of one Tapesty,* R. E. Blahut et al., eds, pages 343–355, Kluwer, Dordrecht and Boston, 1994.
[10] Sason, I.: *On reverse Pinsker inequalities,* Technion, Haifa (2015); `arXiv:1503.07118`.
[11] Verdú, S.: *Total variation distance and the distribution of the relative information,* Proceedings of the 2014 Information Theory and Applications (ITA) Workshop, pp. 499–501, San-Diego, California, USA, February 2014.
[12] Woess, W.: *Denumerable Markov Chains. Generating functions, Boundary Theory, Random Walks on Trees.* European Math. Soc. Publishing House, 2009.
[13] Yari, G. H., and Nikooravesh, Z.: *Estimation of the entropy rate of ergodic Markov chains.* J. Iran. Stat. Soc. (JIRSS) **11** (2012) 75–85.

Institut für Diskrete Mathematik,
Technische Universität Graz,
Steyrergasse 30, A-8010 Graz, Austria

E-mail address: `thirschler@tugraz.at, woess@tugraz.at`